arXiv:1701.03493v1 [cs.DM] 12 Jan 2017

Subgaussian Tail Bounds via Stability Arguments

Thomas Steinke* Jonathan Ullman†

January 16, 2017

Abstract

Sums of independent, bounded random variables concentrate around their expectation approximately as well a Gaussian of the same variance. Well known results of this form include the Bernstein, Hoeffding, and Chernoff inequalities and many others. We present an alternative proof of these tail bounds based on what we call a *stability argument*, which avoids bounding the moment generating function or higher-order moments of the distribution. Our stability argument is inspired by recent work on the generalization properties of differential privacy and their connection to adaptive data analysis (Bassily et al., STOC 2016).

1 Introduction

Tail bounds, also known as large deviation inequalities, or concentration of measure theorems, are an essential tool in mathematics, computer science, and statistics. The most common type of tail bounds (in theoretical computer science) show that the sum of independent, bounded random variables has subgaussian tails—the best known example being the following special case of Bernstein's Inequality.

Theorem 1.1 ([Ber24])**.** *If $X_1, \cdots, X_n$ are independent random variables supported on $[0,1]$ and $\mu_i = \mathbb{E}[X_i]$ for every i, then*

$$\forall \varepsilon \geq 0 \quad \mathbb{P}\left[\sum_{i=1}^{n} X_i - \mu_i \geq \varepsilon n\right] \leq e^{-\Omega(\varepsilon^2 n)}.$$

Theorem 1.1 has many generalizations and extensions: tight constants [Hoe63], random variables with differing ranges [Hoe63], multiplicative error bounds [Che52], variables with low variance [Ber24], variables with only limited independence [SSS95], Martingales [Azu67], matrix-valued random variables [AW02], and many more.

We present an alternative proof of Theorem 1.1, which differs significantly from all other proofs the authors are aware of. Our proof is derived from recent work on algorithmic stability for adaptive data analysis [BNS+16]. At a high level, our proof shows that Theorem 1.1 follows from an appropriate bound on the expectation of the maximum of m independent realizations of $\sum_{i=1}^{n} X_i - \mu_i$. The required bound on the expected maximum follows from the existence of a *stable* randomized algorithm for selecting an approximate maximizer from among the m realizations. Here stability means that the output distribution of the algorithm is insensitive to small changes in these values, which is formalized by differential privacy [DMNS06].

*IBM, Almaden Research Center. dp-tail@thomas-steinke.net

†Northeastern University. jullman@ccs.neu.edu

The alternative proof was (in hindsight) a key ingredient in obtaining tight bounds on the generalization properties of differential privacy [BNS+16], which can be viewed as a strong extension of Theorem 1.1. Thus we believe our approach both deepens our understanding of concentration inequalities and may find further applications.

1.1 Our Approach

Since reasoning directly about the distribution of $\sum_{i=1}^{n} X_i - \mu_i$ is difficult, proofs of Theorem 1.1 typically work by bounding the expectation of some random variable that serves as a "proxy" for the tail probability $\mathbb{P}\left[\sum_{i=1}^{n} X_i - \mu_i \geq \varepsilon n\right]$. Our proof uses a very different proxy than previous proofs. (See Section 1.2 for a comparison to other proofs.) Specifically, to bound $\mathbb{P}[Y \geq y]$ we bound the quantity $\mathbb{E}\left[\max\{0, Y^1, Y^2, \cdots, Y^m\}\right]$, where $Y^1, \cdots, Y^m$ are independent copies of the original random variable Y.[1] Bounding this proxy implies a tail bound via the following lemma.

Lemma 1.2. *Let Y be a random variable and let $Y^1, Y^2, \cdots, Y^m$ be independent copies of Y. Then*

$$\mathbb{P}\left[Y \geq 2\mathbb{E}\left[\max\left\{0, Y^1, \cdots, Y^m\right\}\right]\right] \leq \frac{\ln(2)}{m}.$$

Proof. Let $y = 2\mathbb{E}\left[\max\{0, Y^1, \cdots, Y^m\}\right]$ and $\delta = \mathbb{P}[Y \geq y]$. By Markov's inequality,

$$\mathbb{P}\left[\max\{0, Y^1, Y^2, \cdots, Y^m\} \geq y\right] \leq \frac{1}{2}.$$

However, if $\delta > \ln(2)/m$, then

$$\begin{aligned}\mathbb{P}\left[\max\{0, Y^1, Y^2, \cdots, Y^m\} \geq y\right] &= 1 - \mathbb{P}\left[\forall j \in [m] \quad Y^j < y\right] \\ &= 1 - \mathbb{P}[Y < y]^m = 1 - (1-\delta)^m \\ &> 1 - e^{-\delta m} > 1 - e^{-\ln(2)} = 1/2,\end{aligned}$$

which is a contradiction. Thus $\delta \leq \ln(2)/m$, as required. □

Thus, the key technical step in our proof is to establish the following proxy bound.

Proposition 1.3. *Let $X_1, \cdots, X_n$ be independent random variables supported on $[0,1]$ and $\mu_i = \mathbb{E}[X_i]$ for each i. Define $Y = \sum_{i=1}^{n} X_i - \mu_i$. Then for every $m \in \mathbb{N}$, if $Y^1, \cdots, Y^m$ are independent copies of Y,*

$$\mathbb{E}\left[\max\left\{0, Y^1, \cdots, Y^m\right\}\right] \leq O(\sqrt{n \cdot \ln m}).$$

Combining Proposition 1.3 with Lemma 1.2 and setting $m = e^{\Theta(\varepsilon^2 n)}$ yields Theorem 1.1. We remark that Proposition 1.3 and Theorem 1.1 are equivalent, because the proxy bound can be derived by integrating the tail bound of Theorem 1.1 combined with a union bound.

Our proof of Proposition 1.3 is essentially a special case of the work of [BNS+16] on algorithmic stability, which proves concentration inequalities in a more general "adaptive" setting.

[1] Our proof also works with the proxy $\mathbb{E}\left[\max\{|Y^1|, |Y^2|, \cdots, |Y^m|\}\right]$, which naturally yields two-sided tail bounds.

Bounding the Proxy via Stability

Define $Y^{m+1} = 0$ so that

$$\max\left\{0, Y^1, \cdots, Y^m\right\} = \max\left\{Y^1, \cdots, Y^m, Y^{m+1}\right\} = Y^{\arg\max_{j\in[m+1]} Y^j}.$$

The function $f(Y^1, \cdots, Y^{m+1}) = \arg\max_{j\in[m+1]} Y^j$ is extremely *unstable* in the sense that a small change in the inputs $Y^1, \cdots, Y^{m+1}$ could change the output arbitrarily. In contrast, a constant function is perfectly *stable*, since it doesn't depend on its input at all. Our proof relies on the existence of an intermediate function that approximate the arg max function but provides some stability. We introduce a parameter $\eta > 0$ that parameterizes the degree of stability, and define a procedure $\mathcal{S}_\eta$, and analyze the quantity $\mathbb{E}\left[Y^{\mathcal{S}_\eta(Y^1,\cdots,Y^{m+1})}\right]$.

Smaller values of η imply a higher degree of stability. Specifically, as $\eta \to 0$ then $\mathcal{S}_\eta$ is independent of its input, so $\mathbb{E}\left[Y^{\mathcal{S}_\eta(Y^1,\cdots,Y^{m+1})}\right] \to 0$, since $\mathbb{E}[Y] = 0$. On the other hand, as $\eta \to \infty$, the approximation becomes arbitrarily accurate and we have

$$\mathbb{E}\left[Y^{\mathcal{S}_\eta(Y^1,\cdots,Y^{m+1})}\right] \to \mathbb{E}\left[\max\left\{0, Y^1, \cdots, Y^m\right\}\right].$$

All that remains is to quantify the tradeoff between stability and accuracy, and find a suitable value of $\eta > 0$ that allows us to relate $\mathbb{E}[Y]$ to $\mathbb{E}[\max\{0, Y^1, \cdots, Y^m\}]$. The full proof is contained in Section 2.

1.2 Comparison to Previous Proofs

For context, we give some comparison between our approach and the previous proofs of concentration inequalities. Previous proofs use different proxies for the tails $\mathbb{P}\left[\sum_{i=1}^n X_i - \mu_i \geq \varepsilon n\right]$. The most common proxy is the *moment generating function* $f(\lambda) := \mathbb{E}\left[e^{\lambda \sum_{i=1}^n X_i - \mu_i}\right]$. If $X_1, \cdots, X_n$ are independent random variables ssupported on $[0, 1]$ and $\mu_i = \mathbb{E}[X_i]$ for each i, then

$$\forall \lambda \in \mathbb{R} \quad \mathbb{E}\left[e^{\lambda \sum_{i=1}^n X_i - \mu_i}\right] \leq e^{O(\lambda^2 n)}. \tag{1}$$

By Markov's inequality,[2] (1) implies the tail bound

$$\forall \lambda \in \mathbb{R} \quad \mathbb{P}\left[\sum_{i=1}^n X_i - \mu_i \geq \varepsilon n\right] \leq \frac{\mathbb{E}\left[e^{\lambda(\sum_{i=1}^n X_i - \mu_i)}\right]}{e^{\lambda \varepsilon n}} \leq e^{O(\lambda^2 n) - \lambda \varepsilon n}. \tag{2}$$

Setting $\lambda = \Theta(\varepsilon)$ appropriately in (2) yields Theorem 1.1. The reverse is also true—integrating the tail bound of Theorem 1.1 yields the bound on the moment generating function (1). Thus Theorem 1.1 is equivalent to the proxy bound (1).[3] The moment generating function is particularly easy to work with, as independence can be exploited to "factor" the moment generating function: $\mathbb{E}\left[e^{\lambda \sum_{i=1}^n X_i - \mu_i}\right] = \prod_{i=1}^n \mathbb{E}\left[e^{\lambda(X_i - \mu_i)}\right]$. Thus proving (1) reduces to analysing a single $X_i - \mu_i$ random variable.

[2] Markov's inequality states that, for a non-negative random variable X and a constant $x > 0$, $\mathbb{P}[X \geq x] \leq \mathbb{E}[X]/x$.

[3] More precisely, for *any* centered random variable X, the condition $\forall \lambda \ \ \mathbb{E}\left[e^{\lambda X}\right] \leq e^{O(\lambda^2)}$ is equivalent the the condition $\forall x \geq 0 \ \ \mathbb{P}[|X| \geq x] \leq e^{-\Omega(x^2)}$ (as well as other equivalent characterizations). Such a random variable is called *subgaussian* [Riv12].

The *moments* $m(k) := \mathbb{E}[(\sum_{i=1}^n X_i - \mu_i)^k]$ are another common proxy. If $X_1, \cdots, X_n$ are independent random variables supported on $[0,1]$ and $\mu_i = \mathbb{E}[X_i]$ for each i, then

$$\forall k \in \mathbb{N} \quad \mathbb{E}\left[\left(\sum_{i=1}^n X_i - \mu_i\right)^{2k}\right] \leq O(nk)^k. \tag{3}$$

Theorem 1.1 follows from (3) by applying Markov's inequality with an appropriate choice of $k = \Theta(\varepsilon^2 n)$. Again, the moment bound (3) can be obtained by integrating (a two-sided version of) the tail bound in Theorem 1.1, so moment bounds are equivalent to tail bounds. An immediate benefit of using moment bounds is that the independence assumption in Theorem 1.1 can be relaxed to *limited independence* [SSS95, BS15].

Kabanets and Impagliazzo [IK10] use as a proxy $\mathbb{E}[\prod_{i \in I} X_i]$ where $I \subseteq \{1, 2, \cdots, n\}$ is a random set of indices. Other proofs of Theorem 1.1 tend to resort to a direct analysis of the binomial distribution. We refer the reader to [Mul] for a compendium of proofs.

The Advantages of Different Proxies

The proxy $\mathbb{E}[\max\{0, Y^1, \cdots, Y^m\}]$ is a very natural quantity to bound. In many applications, tail bounds are combined with a union bound because the real quantity of interest is of the form $\max\{Y^1, \cdots, Y^m\}$.[4]

The advantage of our proxy $\mathbb{E}[\max\{0, Y^1, \cdots, Y^m\}]$ is that it less sensitive to heavy tails than either moments or the moment generating function. Namely, suppose that, with probability δ, the random variable Y takes on a large value Δ. Then $\mathbb{E}[\max\{0, Y^1, \cdots, Y^m\}]$ grows linearly with $\Delta \to \infty$ as $O(m\delta\Delta)$, whereas $\mathbb{E}[Y^{2k}]$ grows polynomially as $\delta\Delta^{2k}$ and $\mathbb{E}[e^{\lambda Y}]$ grows exponentially as $O(\delta e^{\lambda\Delta})$. Note that $\mathbb{P}[Y \geq y]$ does not grow at all as Δ increases beyond y.

For Theorem 1.1 the tails are thin enough that all three proxies yield the same result. However, in the application of Bassily et al. [BNS+16] to adaptive data analysis, there is a δ tail event as described above. Thus, using the proxy $\mathbb{E}[\max\{0, Y^1, \cdots, Y^m\}]$ yields tighter bounds than what is achievable using moments. There may be other applications for which our proxy is more suitable than either moments or the moment generating function.

We also note that our proxy can be bounded in terms of the moment generating function: By Jensen's inequality and the bound $\max\{1, e^{\lambda Y^1}, \cdots, e^{\lambda Y^m}\} \leq 1 + e^{\lambda Y^1} + \cdots + e^{\lambda Y^m}$ we can show that

$$\forall \lambda > 0 \quad \mathbb{E}\left[\max\left\{0, Y^1, \cdots, Y^m\right\}\right] \leq \frac{1}{\lambda} \log \mathbb{E}\left[e^{\lambda \cdot \max\{0, Y^1, \cdots, Y^m\}}\right] \leq \frac{1}{\lambda} \log\left(1 + m \cdot \mathbb{E}\left[e^{\lambda Y}\right]\right). \tag{4}$$

Thus Proposition 1.3 can be derived from the moment generating function bound (1) by setting $\lambda = \Theta\left(\sqrt{\ln(m+1)/n}\right)$ appropriately. Likewise, Proposition 1.3 can be derived from the moment bound (3) via

$$\forall k \in \mathbb{N} \quad \mathbb{E}\left[\max\left\{|Y^1|, \cdots, |Y^m|\right\}\right] \leq \mathbb{E}\left[\max\{(Y^1)^{2k}, \cdots, (Y^m)^{2k}\}\right]^{1/2k} \leq \left(m \cdot \mathbb{E}\left[Y^{2k}\right]\right)^{1/2k}. \tag{5}$$

2 Proof of the Proxy Bound

Let $\mathbf{X}$ be a random $n \times m$ matrix with entries supported on $[0,1]$. Let X_i^j denote the random variable corresponding to the entry in the i^{th} row and j^{th} column of $\mathbf{X}$.[5] As a convention we

[4] Proposition 1.3 also holds if the random variables $Y^1, \cdots, Y^m$ are not independent, like the union bound.

[5] We choose the subscript-superscript notation $\mathbf{X}_i^j$, rather than the more common double-subscript notation $\mathbf{X}_{i,j}$, because the rows and columns of $\mathbf{X}$ play markedly different roles that we want to emphasize.

will use lower-case letters $\mathbf{x}$ and $\mathbf{x}_i$ and x_i^j to denote realizations of the matrix, rows of the matrix, and entries of the matrix, respectively, and we will use upper case letters to denote the corresponding random variables.

Lemma 2.1 (Main Lemma)**.** *If $\mathbf{X}$ is a random $n \times m$ matrix with entries supported on $[0,1]$ and independent rows, then*

$$\forall \eta > 0 \qquad \mathbb{E}\left[\max_{j\in[m]} \sum_{i=1}^{n} X_i^j\right] \leq e^{\eta} \max_{j\in[m]} \mathbb{E}\left[\sum_{i=1}^{n} X_i^j\right] + \frac{2\ln(m)}{\eta}.$$

Before delving into the proof, some remarks are in order. Lemma 2.1 states that, under certain independence and boundedness conditions, we can "switch" the expectation and the maximum of a collection of random sums at the price of a multiplicative e^{η} factor and an additive $2\ln(m)/\eta$ factor, for any choice of $\eta > 0$.

In our application, we will start with some independent bounded random variables $(X_1, \ldots, X_n)$. Each column j of $\mathbf{X}$ will an independent copy of these random variables, $(X_1^j, \ldots, X_n^j)$. In this case the columns of $\mathbf{X}$ are also independent. The lemma then says that the the expected maximum of the column sums is not too much larger than the expectation of a single column sum, which implies our tail bounds.

Proof of Lemma 2.1. The proof relies on the existence of a randomized *stable selection procedure* $\mathcal{S}_\eta : [0,1]^{n\times m} \to [m]$ defined by

$$\mathbb{P}\left[\mathcal{S}_\eta(\mathbf{x}) = j\right] = \frac{1}{C_\eta(\mathbf{x})} \exp\left(\tfrac{\eta}{2} \cdot \sum_{i=1}^{n} x_i^j\right), \qquad \text{where} \qquad C_\eta(\mathbf{x}) = \sum_{j=1}^{m} \exp\left(\tfrac{\eta}{2} \cdot \sum_{i=1}^{n} x_i^j\right).$$

The *stability parameter* $\eta > 0$ can be chosen arbitrarily. The key to the proof is that $\mathcal{S}$ balances two goals:

- (Stability) the distribution of $\mathcal{S}_\eta(\mathbf{x})$ is not very sensitive to changing a single row of $\mathbf{x}$ (i.e. the probability mass function can only change by a small multiplicative factor), and
- (Accuracy) in expectation, the chosen column $j = \mathcal{S}_\eta(\mathbf{x})$ is close to the largest column of $\mathbf{x}$ (i.e. $\mathcal{S}_\eta(\mathbf{x}) \approx \arg\max_{j\in[m]} \sum_{i=1}^{n} x_i^j$).

For $\mathbf{x} \in [0,1]^{n\times m}$, $\tilde{\mathbf{x}} \in [0,1]^m$, and $i \in [n]$, define $(\mathbf{x}_{-i}, \tilde{\mathbf{x}}) \in [0,1]^{n\times m}$ to be $\mathbf{x}$ with the i^{th} row replaced by $\tilde{\mathbf{x}}$. That is,

$$(\mathbf{x}_{-i}, \tilde{\mathbf{x}})_{i'}^{j'} = \begin{cases} \mathbf{x}_{i'}^{j'} & \text{if } i' \neq i \\ \tilde{\mathbf{x}}^{j'} & \text{if } i' = i \end{cases}$$

for $i' \in [n]$ and $j' \in [m]$.

The next two claims establish the stability and accuracy properties of $\mathcal{S}_\eta$. These two claims are special cases of well known results in the differential privacy literature.

Claim 2.2 (Stability)**.** *For every $\mathbf{x} \in [0,1]^{n\times m}$, $\tilde{\mathbf{x}} \in [0,1]^m$, $i \in [n]$, and $j \in [m]$,*

$$e^{-\eta}\mathbb{P}\left[\mathcal{S}_\eta(\mathbf{x}) = j\right] \leq \mathbb{P}\left[\mathcal{S}_\eta(\mathbf{x}_{-i}, \tilde{\mathbf{x}}) = j\right] \leq e^{\eta}\mathbb{P}\left[\mathcal{S}_\eta(\mathbf{x}) = j\right].$$

Proof of Claim 2.2. Since

$$\left|\sum_{i'=1}^{n}\mathbf{x}_{i'}^{j}-\sum_{i'=1}^{n}(\mathbf{x}_{-i},\tilde{\mathbf{x}})_{i'}^{j}\right|=\left|\mathbf{x}_{i}^{j}-\tilde{\mathbf{x}}^{j}\right|\leq 1,$$

for every $j\in[m]$, we have

$$e^{-\frac{\eta}{2}}\leq\frac{\exp\left(\frac{\eta}{2}\cdot\sum_{i'=1}^{n}x_{i'}^{j}\right)}{\exp\left(\frac{\eta}{2}\cdot\sum_{i'=1}^{n}(\mathbf{x}_{-i},\tilde{\mathbf{x}})_{i'}^{j}\right)}=\exp\left(\frac{\eta}{2}\cdot\left(\sum_{i'=1}^{n}\mathbf{x}_{i'}^{j}-\sum_{i'=1}^{n}(\mathbf{x}_{-i},\tilde{\mathbf{x}})_{i'}^{j}\right)\right)=\exp\left(\frac{\eta}{2}\cdot\left(\mathbf{x}_{i}^{j}-\tilde{\mathbf{x}}^{j}\right)\right)\leq e^{\frac{\eta}{2}}.$$

Consequently,

$$C_{\eta}(\mathbf{x})=\sum_{j=1}^{m}\exp\left(\frac{\eta}{2}\cdot\sum_{i'=1}^{n}x_{i'}^{j}\right)\leq\sum_{j=1}^{m}e^{\frac{\eta}{2}}\cdot\exp\left(\frac{\eta}{2}\cdot\sum_{i'=1}^{n}(\mathbf{x}_{-i},\tilde{\mathbf{x}})_{i'}^{j}\right)=e^{\frac{\eta}{2}}\cdot C_{\eta}(\mathbf{x}_{-i},\tilde{\mathbf{x}})$$

and, likewise, $C_{\eta}(\mathbf{x}_{-i},\tilde{\mathbf{x}})\leq e^{\frac{\eta}{2}}\cdot C_{\eta}(\mathbf{x})$. The definition of $\mathcal{S}_{\eta}$ implies

$$e^{-\frac{\eta}{2}}\cdot e^{-\frac{\eta}{2}}\leq\frac{\mathbb{P}\left[\mathcal{S}_{\eta}(\mathbf{x}_{-i},\tilde{\mathbf{x}})=j\right]}{\mathbb{P}\left[\mathcal{S}_{\eta}(\mathbf{x})=j\right]}=\frac{\exp\left(\frac{\eta}{2}\cdot\sum_{i'=1}^{n}x_{i'}^{j}\right)}{\exp\left(\frac{\eta}{2}\cdot\sum_{i'=1}^{n}(\mathbf{x}_{-i},\tilde{\mathbf{x}})_{i'}^{j}\right)}\cdot\frac{C_{\eta}(\mathbf{x}_{-i},\tilde{\mathbf{x}})}{C_{\eta}(\mathbf{x})}\leq e^{\frac{\eta}{2}}\cdot e^{\frac{\eta}{2}},$$

as required. □

Claim 2.3 (Accuracy). *For every* $\mathbf{x}\in[0,1]^{n\times m}$,

$$\mathop{\mathbb{E}}_{\mathcal{S}_{\eta}}\left[\sum_{i=1}^{n}x_{i}^{\mathcal{S}_{\eta}(\mathbf{x})}\right]\geq\max_{j\in[m]}\sum_{i=1}^{n}x_{i}^{j}-\frac{2\ln(m)}{\eta}$$

Observe that as $\eta\to\infty$, the accuracy becomes arbitrarily good, whereas the stability degrades.

Proof of Claim 2.3. Recall that, by construction, $\mathcal{S}_{\eta}$ satisfies,

$$\mathbb{P}\left[\mathcal{S}_{\eta}(\mathbf{x})=j\right]=\frac{1}{C_{\eta}(\mathbf{x})}\exp\left(\frac{\eta}{2}\cdot\sum_{i=1}^{n}x_{i}^{j}\right),\qquad\text{where}\qquad C_{\eta}(\mathbf{x})=\sum_{j=1}^{m}\exp\left(\frac{\eta}{2}\cdot\sum_{i=1}^{n}x_{i}^{j}\right).$$

Rearranging, we have

$$\sum_{i=1}^{n}x_{i}^{j}=\frac{2}{\eta}\left(\ln C_{\eta}(\mathbf{x})+\ln\mathbb{P}\left[\mathcal{S}_{\eta}(\mathbf{x})=j\right]\right).$$

Thus,

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathcal{S}_{\eta}}\left[\sum_{i=1}^{n}x_{i}^{\mathcal{S}_{\eta}(\mathbf{x})}\right]=\sum_{j=1}^{m}\mathbb{P}\left[\mathcal{S}_{\eta}(\mathbf{x})=j\right]\cdot\sum_{i=1}^{n}x_{i}^{j}&=\sum_{j=1}^{m}\mathbb{P}\left[\mathcal{S}_{\eta}(\mathbf{x})=j\right]\cdot\frac{2}{\eta}\left(\ln C_{\eta}(\mathbf{x})+\ln\mathbb{P}\left[\mathcal{S}_{\eta}(\mathbf{x})=j\right]\right)\\&=\frac{2}{\eta}\left(1\cdot\ln C_{\eta}(\mathbf{x})+\sum_{j=1}^{m}\mathbb{P}\left[\mathcal{S}_{\eta}(\mathbf{x})=j\right]\ln\mathbb{P}\left[\mathcal{S}_{\eta}(\mathbf{x})=j\right]\right)\\&=\frac{2}{\eta}\left(\ln C_{\eta}(\mathbf{x})-\mathrm{H}\left[\mathcal{S}_{\eta}(\mathbf{x})\right]\right),\end{aligned}$$

where $\mathsf{H}\left[\mathcal{S}_\eta(\mathbf{x})\right]$ is the Shannon entropy of the random variable $\mathcal{S}_\eta(\mathbf{x})$ measured in "nats," rather than bits, as we are using natural logarithms, rather than binary logarithms. Since the random variable $\mathcal{S}_\eta(\mathbf{x})$ is supported on a set of size m, the entropy satisfies $\mathsf{H}\left[\mathcal{S}_\eta(\mathbf{x})\right] \le \ln(m)$, as the entropy function is maximized by the uniform distribution [**?**, Theorem 2.6.4]. By definition we have

$$C_\eta(\mathbf{x}) = \sum_{j=1}^{m} \exp\left(\frac{\eta}{2} \cdot \sum_{i=1}^{n} x_i^j\right) \ge \max_{j\in[m]} \exp\left(\frac{\eta}{2} \sum_{i=1}^{n} x_i^j\right) \quad \text{and} \quad \ln C_\eta(\mathbf{x}) \ge \max_{j\in[m]} \frac{\eta}{2} \cdot \sum_{i=1}^{n} x_i^j.$$

The claim now follows from these two inequalities:

$$\mathop{\mathbb{E}}_{\mathcal{S}_\eta}\left[\sum_{i=1}^{n} x_i^{\mathcal{S}_\eta(\mathbf{x})}\right] = \frac{2}{\eta}\left(\ln C_\eta(\mathbf{x}) - \mathsf{H}\left[\mathcal{S}_\eta(\mathbf{x})\right]\right) \ge \frac{2}{\eta}\left(\max_{j\in[m]} \frac{\eta}{2} \cdot \sum_{i=1}^{n} x_i^j - \ln(m)\right).$$

□

In order to complete the proof of the lemma, we have the following claim. This claim uses stability (Claim 2.2) to show that $\mathop{\mathbb{E}}_{\mathbf{X},\mathcal{S}_\eta}\left[\sum_{i=1}^{n} X_i^{\mathcal{S}_\eta(\mathbf{X})}\right]$ is not much larger than $\mathop{\mathbb{E}}_{\mathbf{X},\mathcal{S}_\eta}\left[\sum_{i=1}^{n} X_i^j\right]$ for a fixed $j \in [m]$.

Claim 2.4. *If* $\mathbb{E}\left[\sum_{i=1}^{n} X_i^j\right] \le \mu$ *for all* $j \in [m]$*, then*

$$\mathop{\mathbb{E}}_{\mathbf{X},\mathcal{S}_\eta}\left[\sum_{i=1}^{n} X_i^{\mathcal{S}_\eta(\mathbf{X})}\right] \le e^\eta \mu.$$

Proof of Claim 2.4. Let $\tilde{\mathbf{X}}$ be an independent and identical copy of $\mathbf{X}$. Let $\tilde{\mathbf{X}}_i \in [0,1]^m$ denote the i^{th} row of $\tilde{\mathbf{X}}$. The key observation is that the pair $(\mathbf{X}_{-i}, \tilde{\mathbf{X}}_i)$ and X_i^j is identically distributed to the pair $\mathbf{X}$ and $\tilde{X}_i^j$ — this is where we use the independence assumption. Thus the pair $\mathcal{S}_\eta(\mathbf{X}_{-i}, \tilde{\mathbf{X}}_i)$ and X_i^j is identically distributed to the pair $\mathcal{S}_\eta(\mathbf{X})$ and $\tilde{X}_i$, which means we can swap them below. By Claim 2.2,

$$\begin{aligned}
\mathop{\mathbb{E}}_{\mathbf{X},\mathcal{S}_\eta}\left[\sum_{i=1}^{n} X_i^{\mathcal{S}_\eta(\mathbf{X})}\right] &= \mathop{\mathbb{E}}_{\mathbf{X}}\left[\sum_{j=1}^{m}\sum_{i=1}^{n} \mathop{\mathbb{P}}_{\mathcal{S}_\eta}\left[\mathcal{S}_\eta(\mathbf{X}) = j\right] X_i^j\right] \\
&\le \mathop{\mathbb{E}}_{\mathbf{X},\tilde{\mathbf{X}}}\left[\sum_{j=1}^{m}\sum_{i=1}^{n} e^\eta \mathop{\mathbb{P}}_{\mathcal{S}_\eta}\left[\mathcal{S}_\eta(\mathbf{X}_{-i}, \tilde{\mathbf{X}}_i) = j\right] X_i^j\right] \\
&= \mathop{\mathbb{E}}_{\mathbf{X},\tilde{\mathbf{X}}}\left[\sum_{j=1}^{m}\sum_{i=1}^{n} e^\eta \mathop{\mathbb{P}}_{\mathcal{S}_\eta}\left[\mathcal{S}_\eta(\mathbf{X}) = j\right] \tilde{X}_i^j\right] \\
&\le \mathop{\mathbb{E}}_{\mathbf{X},\tilde{\mathbf{X}}}\left[\sum_{j=1}^{m} e^\eta \mathop{\mathbb{P}}_{\mathcal{S}_\eta}\left[\mathcal{S}_\eta(\mathbf{X}) = j\right] \mu\right] \\
&= e^\eta \mu.
\end{aligned}$$

□

Combining Claim 2.3 with Claim 2.4 (setting $\mu = \max_{j\in[m]} \mathbb{E}\left[\sum_{i=1}^n \mathbf{X}_i^j\right]$), we have

$$\mathop{\mathbb{E}}_{\mathbf{X}}\left[\max_{j\in[m]} \sum_{i=1}^n X_i^j - \frac{2\ln(m)}{\eta}\right] \le \mathop{\mathbb{E}}_{\mathbf{X},\mathcal{S}_\eta}\left[\sum_{i=1}^n X_i^{\mathcal{S}_\eta(\mathbf{x})}\right] \le e^\eta \mu = e^\eta \max_{j\in[m]} \mathop{\mathbb{E}}_{\mathbf{X}}\left[\sum_{i=1}^n X_i^j\right].$$

Rearranging yields the lemma. □

Lemma 2.1 readily yields the bound claimed in the introduction:

Proposition 2.5 (Proposition 1.3)**.** *Let $X_1, \cdots, X_n$ be independent random variables supported on $[0,1]$ and $\mu_i = \mathbb{E}[X_i]$ for each i. Define $Y = \sum_{i=1}^n X_i - \mu_i$. Fix $m \in \mathbb{N}$ and let $Y^1, \cdots, Y^m$ be independent copies of Y. Then*

$$\mathbb{E}\left[\max\left\{0, Y^1, \cdots, Y^m\right\}\right] \le 4\sqrt{n \cdot \ln(m+1)}.$$

Proof. Firstly, if $m \ge e^n - 1$, then the result holds trivially as $\max\left\{0, Y^1, \cdots, Y^m\right\} \le n$ with certainty. So we may assume $m < e^n - 1$.

Let $\mu = \sum_{i=1}^n \mu_i$. For each $i \in [n]$, let $X_i^1, \cdots, X_i^m$ be independent copies of X_i, so that $Y^j = \sum_{i=1}^n X_i^j - \mu_i$ for all $j \in [m]$. Let $X_i^{m+1} = \mu_i$ be a constant "dummy random varaible" for each i.

Now we apply Lemma 2.1 to the random matrix $\mathbf{X} \in [0,1]^{n\times(m+1)}$:

$$\forall \eta > 0 \quad \mathbb{E}\left[\max_{j\in[m+1]} \sum_{i=1}^n X_i^j\right] \le e^\eta \max_{j\in[m+1]} \mathbb{E}\left[\sum_{i=1}^n X_i^j\right] + \frac{2\ln(m+1)}{\eta}.$$

By construction, $\mathbb{E}\left[\sum_{i=1}^n X_i^j\right] = \sum_{i=1}^n \mu_i = \mu$ for all $j \in [m+1]$. Also $\sum_{i=1}^n X_i^j = Y^j + \mu$ for all $j \in [m]$ and $\sum_{i=1}^n X_i^{m+1} = 0 + \mu$. Substituting in these expressions yields

$$\forall \eta > 0 \quad \mathbb{E}\left[\max\left\{Y^1 + \mu, Y^2 + \mu, \cdots, Y^m + \mu, 0 + \mu\right\}\right] \le e^\eta \mu + \frac{2\ln(m+1)}{\eta}.$$

Subtracting μ gives

$$\forall \eta > 0 \quad \mathbb{E}\left[\max\left\{Y^1, Y^2, \cdots, Y^m, 0\right\}\right] \le (e^\eta - 1)\mu + \frac{2\ln(m+1)}{\eta}.$$

Finally, we use the (crude) bound $\mu \le n$ and the approximation $e^\eta - 1 \le 2\eta$ for $\eta < 1$ to obtain

$$\forall \eta \in (0,1) \quad \mathbb{E}\left[\max\left\{Y^1, Y^2, \cdots, Y^m, 0\right\}\right] \le 2\eta n + \frac{2\ln(m+1)}{\eta}.$$

Set $\eta = \sqrt{\ln(m+1)/n} < \sqrt{\ln(e^n)/n} = 1$ to complete the proof. □

Now we can prove the subgaussian tail bound using the proxy bound of Proposition 2.5.

Theorem 2.6 (Theorem 1.1)**.** *If $X_1, \cdots, X_n$ are independent random variables supported on $[0,1]$ and $\mu_i = \mathbb{E}[X_i]$ for every i, then*

$$\forall \varepsilon \ge 0 \quad \mathbb{P}\left[\sum_{i=1}^n X_i - \mu_i \ge \varepsilon n\right] \le e^{1-\varepsilon^2 n/64}.$$

Needless to say, we have not optimized the constants in Theorem 2.6.

Proof. Fix $m \in \mathbb{N}$ to be determined later. Let $Y = \sum_{i=1}^n X_i - \mu$ and let $Y^1, \cdots, Y^m$ be independent copies of Y. By Proposition 2.5, we have

$$\mathbb{E}\left[\max\{0, Y^1, \cdots, Y^m\}\right] \le 4\sqrt{n\ln(m+1)}.$$

By Lemma 1.2, $\mathbb{P}\left[Y \ge 2\mathbb{E}\left[\max\{0, Y^1, \cdots, Y^m\}\right]\right] \le \ln(2)/m$. Now, set $m = \left\lfloor e^{\varepsilon^2 n/64} - 1 \right\rfloor$, so that $8\sqrt{n\ln(m+1)} \le \varepsilon n$. Then

$$\mathbb{P}[Y \ge \varepsilon n] \le \mathbb{P}\left[Y \ge 8\sqrt{n\ln(m+1)}\right] \le \mathbb{P}\left[Y \ge 2\mathbb{E}\left[\max\{0, Y^1, \cdots, Y^m\}\right]\right] \le \frac{\ln(2)}{m} \le \frac{\ln(2)}{e^{\varepsilon^2 n/64} - 2}.$$

Thus

$$\mathbb{P}[Y \ge \varepsilon n] \le \min\left\{1, \frac{\ln(2)}{e^{\varepsilon^2 n/64} - 2}\right\} \le (2 + \ln 2) \cdot e^{-\varepsilon^2 n/64} \le e^{1 - \varepsilon^2 n/64}.$$

□

3 Extensions

To demonstrate the versatility of our techniques, we show how they can be used to prove some well-known generalizations of Theorem 1.1.

3.1 Multiplicative Tail Bound

Theorem 3.1. *If $X_1, \ldots, X_n$ are independent random variables supported on $[0,1]$ and $\mu = \mathbb{E}\left[\sum_{i=1}^n X_i\right]$, then*

$$\forall \varepsilon \ge 0 \qquad \mathbb{P}\left[\sum_{i=1}^n X_i \ge (1+\varepsilon)\mu\right] \le e \cdot \left(1 + \frac{\varepsilon}{4}\right)^{-\varepsilon\mu/8}.$$

Moreover, the above bound implies

$$\forall \varepsilon \in [0, 10] \qquad \mathbb{P}\left[\sum_{i=1}^n X_i \ge (1+\varepsilon)\mu\right] \le e^{1 - \varepsilon^2\mu/64}.$$

Proof. Let $\delta = \mathbb{P}\left[\sum_{i=1}^n X_i \ge (1+\varepsilon)\mu\right]$ and $m = \lceil \ln(2)/\delta \rceil$. For each $i \in [n]$, let $X_i^1, \cdots, X_i^m$ be independent copies of X_i and let $X_i^{m+1} = \mathbb{E}[X]$ be a constant. Now we apply Lemma 2.1 to the random matrix $\mathbf{X} \in [0,1]^{n \times (m+1)}$ to obtain

$$\mathbb{E}\left[\max_{j \in [m+1]} \sum_{i=1}^n X_i^j\right] \le e^\eta \mu + \frac{2\ln(m+1)}{\eta}. \tag{6}$$

On the other hand,

$$\mathbb{E}\left[\max_{j \in [m+1]} \sum_{i=1}^n X_i^j\right] \ge \mu + \varepsilon\mu \cdot \mathbb{P}\left[\max_{j \in [m]} \sum_{i=1}^n X_i^j \ge (1+\varepsilon)\mu\right] = \mu + \varepsilon\mu(1 - (1-\delta)^m) \ge \mu + \varepsilon\mu(1 - e^{-\delta m}) \ge \mu\left(1 + \frac{\varepsilon}{2}\right).$$

Combining this with (6) and subtracting μ from both sides gives

$$\forall \eta > 0 \qquad \frac{\varepsilon}{2}\mu \le (e^{\eta} - 1)\,\mu + \frac{2\ln(m+1)}{\eta}.$$

Now we set $\eta = \ln(1+\varepsilon/4)$ and rearrange to get $\ln(m+1) \ge \frac{\varepsilon}{8}\mu \ln\left(1+\frac{\varepsilon}{4}\right)$. Since $m \le \ln(2)/\delta + 1$, we have

$$\mathbb{P}\left[\sum_{i=1}^{n} X_i \ge (1+\varepsilon)\mu\right] = \delta \le \min\left\{1, \frac{\ln 2}{m-1}\right\} \le \frac{2+\ln 2}{m+1} \le e\left(1+\frac{\varepsilon}{4}\right)^{-\varepsilon\mu/8}.$$

This gives the first half of the theorem. The second half follows from the fact that, if $0 \le \varepsilon \le 10$, then $1 + \varepsilon/4 \ge \exp(\varepsilon/8)$. □

We can also bound $\mathbb{P}\left[\sum_{i=1}^{n} X_i \le (1-\varepsilon)\mu\right]$. This requires changing Lemma 2.1 to the following.

Lemma 3.2. *If* $\mathbf{X}$ *is a random* $n \times m$ *matrix with entries supported on* $[0,1]$ *and independent rows, then*

$$\forall \eta > 0 \quad \mathbb{E}\left[\min_{j\in[m]} \sum_{i=1}^{n} X_i^j\right] \ge e^{-\eta} \min_{j\in[m]} \mathbb{E}\left[\sum_{i=1}^{n} X_i^j\right] - \frac{2\ln(m)}{\eta}.$$

The proof of Lemma 3.2 is almost identical to that of Lemma 2.1, except the maximum is replaced with the minimum, η is replaced with $-\eta$, and some inequalities are reversed.

Lemma 3.2 yields the following bound.

Theorem 3.3. *If* $X_1, \ldots, X_n$ *are independent random variables supported on* $[0,1]$ *and* $\mu = \mathbb{E}\left[\sum_{i=1}^{n} X_i\right]$, *then*

$$\forall \varepsilon \in [0,1] \quad \mathbb{P}\left[\sum_{i=1}^{n} X_i \le (1-\varepsilon)\mu\right] \le e^{1-\varepsilon^2\mu/32}.$$

3.2 McDiarmid's Inequality

Subgaussian tail bounds are not specific to summation. For independent random variables $X_1, \cdots, X_n$, we can prove subgaussian tail bounds on $f(X_1, \cdots, X_n)$ for any "low-sensitivity" function f. The sum — that is, $f(x_1, \cdots, x_n) = \sum_{i=1}^{n} x_i$ — is only one example of a low-sensitivity function. The more general property that we require is below.

Definition 3.4. A function $f : \mathcal{X}^n \to \mathbb{R}$ is sensitivity-Δ if, for all $i \in [n]$ and all $x_1, x_2, \cdots, x_n, x_i' \in \mathcal{X}$, we have

$$\left|f(x_1, x_2, \cdots, x_n) - f(x_1, x_2, \cdots, x_{i-1}, x_i', x_{i+1}, \cdots, x_n)\right| \le \Delta.$$

We can generalize Proposition 1.3 to low-sensitivity functions as follows.

Lemma 3.5. *Let* $\mathbf{X}$ *be a random* $n \times m$ *matrix with the entries supported on* $\mathcal{X}$ *(not necessarily being real numbers). Assume that the rows of* $\mathbf{X}$ *are independent. Let* $f : \mathcal{X}^n \to \mathbb{R}$ *be sensitivity-*Δ. *Then*

$$\mathbb{E}\left[\max_{j\in[m]}\left(f(X_1^j, \cdots, X_n^j) - \mathbb{E}\left[f(X_1^j, \cdots, X_n^j)\right]\right)\right] \le 8\Delta\sqrt{n \ln m}.$$

The proof of Lemma 3.5 is similar to that of Theorem 7.2 of Bassily et al. [BNS+16]. Lemma 3.5 allows us to generalize Theorem 1.1 to McDiarmid's inequality:

Theorem 3.6 ([McD89])**.** *Let* $X_1, \cdots, X_n$ *be independent random variables and let* $f : \mathcal{X}^n \to \mathbb{R}$ *have sensitivity-*Δ. *Then, for all* $\varepsilon \ge 0$,

$$\mathbb{P}\left[f(X_1, \cdots, X_n) - \mathbb{E}\left[f(X_1, \cdots, X_n)\right] \ge \varepsilon n \Delta\right] \le e^{-\Omega(\varepsilon^2 n)}.$$

References

[AW02] Rudolf Ahlswede and Andreas Winter. Strong converse for identification via quantum channels. *IEEE Transactions on Information Theory*, 48(3):569–579, 2002.

[Azu67] Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal, Second Series*, 19(3):357–367, 1967.

[Ber24] Sergei Bernstein. On a modification of chebyshev?s inequality and of the error formula of laplace. *Ann. Sci. Inst. Sav. Ukraine, Sect. Math*, 1(4):38–49, 1924.

[BNS+16] Raef Bassily, Kobbi Nissim, Adam Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1046–1059. ACM, 2016.

[BS15] Mark Bun and Thomas Steinke. Weighted polynomial approximations: Limits for learning and pseudorandomness. In *RANDOM*, volume 19, pages 625–644. Leibniz International Proceedings in Informatics, 2015.

[Che52] Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, pages 493–507, 1952.

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, 2006.

[Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.

[IK10] Russell Impagliazzo and Valentine Kabanets. Constructive proofs of concentration bounds. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 617–631. Springer, 2010.

[McD89] Colin McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989.

[Mul] Wolfgang Mulzer. Chernoff bounds. https://page.mi.fu-berlin.de/mulzer/notes/misc/chernoff.pdf. Accessed 16-Dec-2016.

[Riv12] Omar Rivasplata. Subgaussian random variables: An expository note. https://sites.ualberta.ca/~omarr/publications/subgaussians.pdf, 2012.

[SSS95] Jeanette P Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff-hoeffding bounds for applications with limited independence. *SIAM Journal on Discrete Mathematics*, 8(2):223–250, 1995.